EAST Search History

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	25917008	@ad<"20040422"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L2	0	L1 and search$3 with segment with sequence same cache same way	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L3	1234	L1 and search$3 with (segment way) with sequence	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L4	12	L1 and search$3 with way with sequence same cycle	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L5	19	L1 and search$3 with segment with sequence same cycle	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L6	0	L1 and search$3 with sequence same way same segment same cycle	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L7	4	L1 and search$3 same sequence same way same segment same cycle	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14

EAST Search History

L8	2	"5778427".pn.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L9	2	"5497477".pn.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L10	2	"6748492".pn.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:14
L11	2	1 and invalidated adj (way block) and (segment line) same (way block) same search$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:22
L12	0	1 and (segment line) same (way block) same search$3 same "not" near3 accessed	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:26
L13	0	1 and (segment line) same (way block) same search$3 same accessed near3 "yet"	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:38
L14	51	1 and (segment line) same (way block) same search$3 same accessed same (cycle sequence)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:49

EAST Search History

L15	21	1 and (segment line) same (way block) same search$3 same accessed same (cycle sequence) same cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:33
L16	267	steely.in.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:53
L17	1	16 and "cache way".clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:54
L18	1	16 and "way search".clm.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:55
L19	43	steely.in.	US-PGPUB	OR	ON	2007/03/30 12:55
L20	1	19 and "way search".clm.	US-PGPUB	OR	ON	2007/03/30 12:55
L21	1171	711/133.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:56
L22	254	711/134.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:56
L23	418	711/135.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:56

EAST Search History

L24	515	711/136.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:56
L25	2152	711/118.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:56
L26	1827	365/49.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 12:57
L27	270	365/50.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:09
L28	2	"20050240731".pn.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:09
L29	693	711/128.ccls.	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:28
L30	19	(21 22 23 24 25 26 27 29) and (3 4 5 7 11 15)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:54

EAST Search History

L31	6	1 and (segment line) same (way block) same search$3 same accessed same (cycle sequence) and associative with cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:38
L32	10	1 and (segment line page) same (way block) same search$3 same accessed same (cycle sequence) and associative with cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:45
L33	73	1 and (segment line page) same (way block) same search$3 same (cycle sequence) and associative with cache	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:46
L34	7	33 and replacement adj polic$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:47
L35	0	33 and "not" near3 accessed	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:49
L36	21	33 and accessed with search$3	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:50
L37	48	(21 22 23 24 25 26 27 29) and (3 4 5 7 11 14 15 31 32 33 34 36)	US-PGPUB; USPAT; USOCR; FPRS; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2007/03/30 13:55

☐ Search Session History

Fri, 30 Mar 2007, 3:18:00 PM EST

Edit an existing query or compose a new query in the Search Query Display.

Select a search number (#) to:

- Add a query to the Search Query Display
- Combine search queries using AND, OR, or NOT
- Delete a search
- Run a search

Search Query Display

[redacted]

Recent Search Queries

#1 ((associative cache<in>metadata) <and> (way search<in>metadata))<and> (segment search<in>metadata)

#2 ((ways<in>metadata) <and> (segments<in>metadata)) <and> (search<in>metadata)

#3 ((cache<in>metadata) <and> (way search<in>metadata)) <and> (segments<in>metadata)

#4 ((cache set<in>metadata) <and> (way<in>metadata))<and> (segment<in>metadata)

#5 ((accessed<in>metadata) <and> (searching<in>metadata)) <and> (way<in>metadata)

#6 ((way search<in>metadata) <and> (search cycle<in>metadata))<and> (segment<in>metadata)

#7 ((cache set<in>metadata) <and> (search<in>metadata)) <and> (accessed<in>metadata)

[redacted]

Terms used
cache set search sequence cycle way accessed segment line block

Found **1,709** of **198,991**

Sort results by: relevance
Display results: expanded form

Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn
August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**
Publisher: ACM Press
Full text available: pdf(63.03 MB) Additional Information: full citation, abstract, citings

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

2 External memory algorithms and data structures: dealing with **massive data**

Jeffrey Scott Vitter
June 2001 **ACM Computing Surveys (CSUR)**, Volume 33 Issue 2
Publisher: ACM Press
Full text available: pdf(828.46 KB) Additional Information: full citation, abstract, references, citings, index terms

Data sets in large applications are often too massive to fit completely inside the computers internal memory. The resulting input/output communication (or I/O) between fast internal memory and slower external memory (such as disks) can be a major performance bottleneck. In this article we survey the state of the art in the design and analysis of external memory (or EM) algorithms and data structures, where the goal is to exploit locality in order to reduce the I/O costs. We consider a varie ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, hierarchical memory, multidimensional access methods, multilevel memory, online, out-of-core, secondary storage, sorting

3 Cache Memories

Alan Jay Smith
September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3
Publisher: ACM Press

Full text available: pdf(4.61 MB) Additional Information: full citation, references, citings, index terms

4 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

5 Energy-performance trade-offs for spatial access methods on memory-resident data

Ning An, Sudhanva Gurumurthi, Anand Sivasubramaniam, Narayanan Vijaykrishnan, Mahmut Kandemir, Mary Jane Irwin

November 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(641.55 KB) Additional Information: full citation, abstract, citings, index terms

The proliferation of mobile and pervasive computing devices has brought energy constraints into the limelight. Energy-conscious design is important at all levels of system architecture, and the software has a key role to play in conserving battery energy on these devices. With the increasing popularity of spatial database applications, and their anticipated deployment on mobile devices (such as road atlases and GPS-based applications), it is critical to examine the energy implications of spatial ...

Keywords: Energy optimization, Multidimensional indexing, Resource-constrained computing, Spatial data

6 Improving instruction cache performance in OLTP

Stavros Harizopoulos, Anastassia Ailamaki

September 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31 Issue 3

Publisher: ACM Press

Full text available: pdf(783.16 KB) Additional Information: full citation, abstract, references, index terms

Instruction-cache misses account for up to 40% of execution time in online transaction processing (OLTP) database workloads. In contrast to data cache misses, instruction misses cannot be overlapped with out-of-order execution. Chip design limitations do not allow increases in the size or associativity of instruction caches that would help reduce misses. On the contrary, the effective instruction cache size is expected to further decrease with the adoption of multicore and multithreading ...

Keywords: Instruction cache, cache misses

7 Memory-wall: Boosting trace cache performance with nonhead miss speculation

Stevan Vlaovic, Edward S. Davidson

June 2002 **Proceedings of the 16th international conference on Supercomputing ICS '02**

Publisher: ACM Press

Full text available: pdf(179.52 KB) Additional Information: full citation, abstract, references, index terms

Trace caches are used to help dynamic branch prediction make multiple predictions in a cycle by embedding some of the predictions in the trace. In this work, we evaluate a trace cache that is capable of delivering a trace consisting of a variable number of instructions

via a linked list mechanism. We evaluate several schemes in the context of an x86 processor model that stores decoded instructions. By developing a new classification for trace cache accesses, we are able to target those misses t ...

Keywords: branch prediction, optimization, trace cache, x86

8 Cache Refill/Access Decoupling for Vector Machines

Christopher Batten, Ronny Krashinsky, Steve Gerding, Krste Asanovic

December 2004 **Proceedings of the 37th annual IEEE/ACM International Symposium on Microarchitecture MICRO 37**

Publisher: IEEE Computer Society

Full text available: pdf(319.32 KB) Additional Information: full citation, abstract, citings

Vector processors often use a cache to exploit temporal locality and reduce memory bandwidth demands, but then require expensive logic to track large numbers of outstanding cache misses to sustain peak bandwidth from memory. We present refill/access decoupling, which augments the vector processor with a Vector Refill Unit (VRU) to quickly pre-execute vector memory commands and issue any needed cache line refills ahead of regular execution. The VRU reduces costs by eliminating much of the outstan ...

9 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research CASCON '97**

Publisher: IBM Press

Full text available: pdf(4.21 MB) Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

10 Inverted files for text search engines

Justin Zobel, Alistair Moffat

July 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 2

Publisher: ACM Press

Full text available: pdf(944.29 KB) Additional Information: full citation, abstract, references, index terms

The technology underlying text search engines has advanced dramatically in the past decade. The development of a family of new index representations has led to a wide range of innovations in index storage, index construction, and query evaluation. While some of these developments have been consolidated in textbooks, many specific techniques are not widely known or the textbook descriptions are out of date. In this tutorial, we introduce the key techniques in the area, describing both a core impl ...

Keywords: Inverted file indexing, Web search engine, document database, information retrieval, text retrieval

11 System-level power optimization: techniques and tools

Luca Benini, Giovanni de Micheli

April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2

Publisher: ACM Press

Full text available: pdf(385.22 KB) Additional Information: full citation, abstract, references, citings, index terms

This tutorial surveys design methods for energy-efficient system-level design. We consider electronic sytems consisting of a hardware platform and software layers. We consider the three major constituents of hardware that consume energy, namely computation, communication, and storage units, and we review methods of reducing their energy consumption. We also study models for analyzing the energy cost of software, and methods for energy-efficient software design and compilation. This survery ...

12 Exploiting perception in high-fidelity virtual environments: Exploiting perception in high-fidelity virtual environments

Additional presentations from the 24th course are available on the citation page

Mashhuda Glencross, Alan G. Chalmers, Ming C. Lin, Miguel A. Otaduy, Diego Gutierrez
July 2006 **ACM SIGGRAPH 2006 Courses SIGGRAPH '06**

Publisher: ACM Press

Full text available: pdf(5.07 MB) mov(68:6 MIN) Additional Information: full citation, abstract, references

The objective of this course is to provide an introduction to the issues that must be considered when building high-fidelity 3D engaging shared virtual environments. The principles of human perception guide important development of algorithms and techniques in collaboration, graphical, auditory, and haptic rendering. We aim to show how human perception is exploited to achieve realism in high fidelity environments within the constraints of available finite computational resources.In this course w ...

Keywords: collaborative environments, haptics, high-fidelity rendering, human-computer interaction, multi-user, networked applications, perception, virtual reality

13 Optimizing instruction cache performance of embedded systems

S. Bartolini, C. A. Prete
November 2005 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 4 Issue 4

Publisher: ACM Press

Full text available: pdf(817.74 KB) Additional Information: full citation, abstract, references, index terms, review

In the embedded domain, the gap between memory and processor performance and the increase in application complexity need to be supported without wasting precious system resources: die size, power, etc. For these reasons, effective exploitation of small and simple cache memories is of the utmost importance. However, programs running on such caches can experience serious inefficiencies due to cache conflicts.We present a new Cache-Aware Code Allocation Technique (CAT), which transforms the structu ...

Keywords: Embedded systems, cache performance, code generation, code reordering, conflict miss

14 Improving the energy behavior of block buffering using compiler optimizations

M. Kandemir, J. Ramanujam, U. Sezer
January 2006 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 11 Issue 1

Publisher: ACM Press

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, index terms

On-chip caches consume a significant fraction of the energy in current microprocessors. As a result, architectural/circuit-level techniques such as block buffering and sub-banking have been proposed and shown to be very effective in reducing the energy consumption of on-chip caches. While there has been some work on evaluating the energy and performance impact of different block buffering schemes, we are not aware of software solutions to take advantage of on-chip cache block buffers.This articl ...

Keywords: Energy optimizations, block buffering, compiler transformations, data cache, embedded systems

15 RISCY patents

David A. Patterson
September 1988 **ACM SIGARCH Computer Architecture News**, Volume 16 Issue 4
Publisher: ACM Press
Full text available: pdf(1.83 MB) Additional Information: full citation, index terms

16 4.2BSD and 4.3BSD as examples of the UNIX system

John S. Quarterman, Abraham Silberschatz, James L. Peterson
December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4
Publisher: ACM Press
Full text available: pdf(4.07 MB) Additional Information: full citation, abstract, references, citings, index terms, review

This paper presents an in-depth examination of the 4.2 Berkeley Software Distribution, Virtual VAX-11 Version (4.2BSD), which is a version of the UNIX Time-Sharing System. There are notes throughout on 4.3BSD, the forthcoming system from the University of California at Berkeley. We trace the historical development of the UNIX system from its conception in 1969 until today, and describe the design principles that have guided this development. We then present the internal data structures and ...

17 Power reduction techniques for microprocessor systems

Vasanth Venkatachalam, Michael Franz
September 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 3
Publisher: ACM Press
Full text available: pdf(602.33 KB) Additional Information: full citation, abstract, references, index terms

Power consumption is a major factor that limits the performance of computers. We survey the "state of the art" in techniques that reduce the total power consumed by a microprocessor system over time. These techniques are applied at various levels ranging from circuits to architectures, architectures to system software, and system software to applications. They also include holistic approaches that will become more important over the next decade. We conclude that power management is a ...

Keywords: Energy dissipation, power reduction

18 1 - Regular Papers: Timing analysis for preemptive multitasking real-time systems with caches

Yudong Tan, Vincent Mooney
February 2007 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 6 Issue 1
Publisher: ACM Press
Full text available: pdf(398.75 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we propose an approach to estimate the worst-case response time (WCRT) of each task in a preemptive multitasking single-processor real-time system utilizing an L1 cache. The approach combines intertask cache-eviction analysis and intratask cache-access analysis to estimate the number of cache lines that can possibly be evicted by the preempting task and also be accessed again by the preempted task after preemptions (thus requiring the preempted task to reload the cache line(s)). T ...

Keywords: Real-time, worst-case response time

19 Balanced Cache: Reducing Conflict Misses of Direct-Mapped Caches

Chuanjun Zhang

May 2006 **ACM SIGARCH Computer Architecture News , Proceedings of the 33rd annual international symposium on Computer Architecture ISCA '06**, Volume 34 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(391.30 KB) Additional Information: full citation, abstract, index terms

Level one cache normally resides on a processor's critical path, which determines the clock frequency. Directmapped caches exhibit fast access time but poor hit rates compared with same sized set-associative caches due to nonuniform accesses to the cache sets, which generate more conflict misses in some sets while other sets are underutilized. We propose a technique to reduce the miss rate of direct mapped caches through balancing the accesses to cache sets. We increase the decoder length and th ...

20 Real-time volume graphics

Klaus Engel, Markus Hadwiger, Joe M. Kniss, Aaron E. Lefohn, Christof Rezk Salama, Daniel Weiskopf

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available: pdf(7.63 MB) Additional Information: full citation, abstract

The tremendous evolution of programmable graphics hardware has made high-quality real-time volume graphics a reality. In addition to the traditional application of rendering volume data in scientific visualization, the interest in applying these techniques for real-time rendering of atmospheric phenomena and participating media such as fire, smoke, and clouds is growing rapidly. This course covers both applications in scientific visualization, e.g., medical volume data, and real-time rendering, ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next